



US010846001B2

(12) **United States Patent**  
**Vallala et al.**

(10) **Patent No.:** **US 10,846,001 B2**

(45) **Date of Patent:** **Nov. 24, 2020**

(54) **ALLOCATING STORAGE REQUIREMENTS  
IN A DISTRIBUTED STORAGE SYSTEM**

(71) Applicant: **Robin Systems, Inc.**, San Jose, CA  
(US)

(72) Inventors: **Shravan Kumar Vallala**, San Jose, CA  
(US); **Dhanashankar Venkatesan**, San  
Jose, CA (US)

(73) Assignee: **ROBIN SYSTEMS, INC.**, San Jose,  
CA (US)

(\*) Notice: Subject to any disclaimer, the term of this  
patent is extended or adjusted under 35  
U.S.C. 154(b) by 0 days.

(21) Appl. No.: **15/807,035**

(22) Filed: **Nov. 8, 2017**

(65) **Prior Publication Data**

US 2019/0138235 A1 May 9, 2019

(51) **Int. Cl.**  
**G06F 3/06** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **G06F 3/0631** (2013.01); **G06F 3/061**  
(2013.01); **G06F 3/0604** (2013.01); **G06F**  
**3/067** (2013.01); **G06F 3/0659** (2013.01);  
**G06F 3/0665** (2013.01)

(58) **Field of Classification Search**  
CPC .... G06F 3/0631; G06F 3/0604; G06F 3/0665;  
G06F 3/067  
See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

4,310,883 A \* 1/1982 Clifton ..... G06F 9/5016  
5,602,993 A 2/1997 Stromberg

6,014,669 A 1/2000 Slaughter  
6,052,797 A 4/2000 Ofek  
6,119,214 A 9/2000 Dirks  
6,157,963 A 12/2000 Courtright, II  
6,161,191 A 12/2000 Slaughter  
6,298,478 B1 10/2001 Nally  
6,301,707 B1 10/2001 Carroll  
6,311,193 B1 10/2001 Sekido  
6,851,034 B2 2/2005 Challenger  
6,886,160 B1 4/2005 Lee

(Continued)

FOREIGN PATENT DOCUMENTS

WO WO2017008675 1/2017

OTHER PUBLICATIONS

Segment map, Feb. 2019.

(Continued)

*Primary Examiner* — William B Partridge

*Assistant Examiner* — Eric T Loonan

(74) *Attorney, Agent, or Firm* — David R. Stevens;  
Stevens Law Group

(57) **ABSTRACT**

A distributed storage schemes manages implementation of QoS targets for IOPs across compute nodes executing applications, primary storage nodes storing a primary copy of a logical storage volume, and clone storage nodes. On the compute node, a maximum priority is assigned to a minimum number of IOPs in a queue within a time window from a time of receipt of a last unexecuted IOP. Other IOPs are assigned a minimum priority. On the storage node, maximum priority IOPs are assigned to high priority queues, from which IOPs are executed first, and low priority IOPs are assigned to low priority queues. Methods for determining the capacity of storage nodes and allocating storage requests are also disclosed.

**20 Claims, 9 Drawing Sheets**

